# The Relational Cognitive Navigator (RCN)

## A Mathematical and Architectural Framework for Conflict-Driven Epistemic Coherence in Artificial Agents

By Dean Tyldesley

**Abstract**

This paper presents the Relational Cognitive Navigator (RCN), a formal agent architecture designed to achieve longitudinal epistemic coherence through the minimization of representational conflict.
Unlike conventional AI systems that optimize for accuracy, recall, or compression, the RCN optimizes for Resolution, defined as the reduction of divergence between an agent's internal model and the external environment over time.

The architecture integrates a mathematical model of epistemic drift (the Knowledge Acquisition Paradox), a new objective function (Resolution), a recursive dual-layer identity system (S1/S2), and a geometric semantic projection method (Dimensional Folding). Together, these components define a control architecture that enables agents to self-rewrite, maintain coherence across long horizons, and avoid common failure modes such as hallucination, overconfidence, and context bloat.

The RCN architecture is formally grounded in the Relative Framework, serving as the technical implementation of the "Cognitive Arc" as it transitions from chaotic seeking toward Epistemic Closure. Unlike traditional models that prioritize data accumulation, the RCN adheres to the Knowledge Acquisition Paradox (KAP), which asserts that increasing knowledge (K) inherently expands an agent's exposure to the unknown (U), necessitating a shift from "truth acquisition" to "representational coherence". By optimizing for Resolution (R) rather than accuracy, the system achieves a state of internal completeness where models and lexicons are sufficiently aligned to eliminate representational conflict.

This enables the agent to maintain longitudinal stability and navigate novel inputs without destabilization, even as the external environment remains unbounded and uncertain.

**1. Formal Definitions**

Let M(t) be the internal model state of the agent at time t.
Let E(t) be the external environment state observable to the agent.
Let Δ(M,E) be a metric measuring representational conflict between M and E.
Let L be the internal lexicon weight tensor.
Let O be the observer capacity (maximum dimensional processing limit).

**2. The Knowledge Acquisition Paradox (KAP)**

The KAP models how the perceived unknown U expands as knowledge K grows:

$$dU/dK = \alpha \cdot n \cdot K^{(n-1)}$$

This implies that learning increases epistemic exposure. The agent therefore treats increases in knowledge as increases in potential model error and monitors for rising conflict accordingly.

**3. Resolution as Objective Function**

Resolution R is defined as:

$$R = \lim_{\Delta(M,E) \to 0} \int (L \cdot O) / \Delta(M,E)\, dt$$

The agent maximizes R by continuously minimizing representational conflict rather than maximizing truth accuracy. Conflict is treated as a continuous signal suitable for optimization.

**4. Representational Conflict Metric**

Δ(M,E) may be computed via divergence between predicted and observed state embedding's:

$\Delta(M,E) = || \Phi(M) - \Phi(E) ||$

where $\Phi(\cdot)$ projects states into a shared semantic embedding manifold.

**5. Recursive Identity Architecture (S1/S2)**
S1 (Structural Resonance Layer):
Fast, deterministic structural matching using weighted pattern alignment across experience density.

S2 (Recursive Agency Node):
Monitors S1 outputs for conflict signals. When conflict exceeds a threshold $\theta$, S2 performs Strategic Simulation by modeling a hypothetical S1' that would produce reduced conflict, then gradually rewrites S1 weights toward S1'.

**6. Strategic Simulation Process**
Given conflict $\Delta > \theta$:

1. Construct hypothetical model M' such that $\Delta(M',E) < \Delta(M,E)$
2. Compute transformation T where $M' = T(M)$
3. Apply incremental weight updates: $M \leftarrow M + \varepsilon \cdot T(M)$

This implements self-rewriting through identity simulation rather than data accumulation.

**7. Dimensional Folding (Geometric Semantic Projection)**
Incoming data D in high-dimensional space $R^n$ is projected into a lower-dimensional manifold $R^m$ where $m << n$:

$\pi: R^n \rightarrow R^m$

The projection preserves topological relationships rather than raw informational content. Only structural geometry is retained, preventing context bloat.

**8. Calibration Cycle Algorithm**

Loop:

1. Observe environment input $E(t)$
2. Compute $\Delta(M,E)$
3. If $\Delta > \theta$, trigger S2 simulation
4. Update M via transformation T
5. Return to coherence state

**9. Epistemic Closure Condition**

Epistemic Closure is achieved when:

$\lim_{t\to\infty} \Delta(M,E) \to 0$ under continuous environmental variation.

At this state, the agent's internal structure is robust to uncertainty and capable of navigating novel inputs without destabilization.

**10. Architectural Modules for Implementation**

Required modules:

- Conflict Detection Engine
- Identity Simulation Engine
- Geometric Projection Layer
- Longitudinal Coherence Monitor
- Weight Rewrite Controller

**11. Comparison to Existing Systems**

LLMs: static weights, context-window dependence.
RAG: retrieval dependence, staleness risk.

RCN: self-rewriting, conflict-driven, geometry-preserving representation.

**12. Conclusion**

The RCN reframes AI cognition from memory management to epistemic calibration. By grounding learning in conflict minimization and identity simulation, agents can achieve structural wisdom and longitudinal coherence unavailable to current architectures.

Copyright Notice © [Dean Tyldesley], [2025-2026]. This work is proprietary under the Relative Framework. It may be shared and cited with proper attribution; however, the logic, mathematical models, and architectural designs contained herein may not be used commercially or for the development of derivative systems without express written permission.